

Occam's Vorpall Quantum Razor: Memory Reduction When Simulating Continuous-Time Stochastic Processes With Quantum Devices

Thomas J. Elliott^{1,*} and Mile Gu^{1,2,3,†}

¹*School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 639673*

²*Complexity Institute, Nanyang Technological University, Singapore 639673*

³*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*

(Dated: April 14, 2017)

Continuous-time stochastic processes pervade everyday experience, and the simulation of models of these processes is of great utility. Classical models of systems operating in continuous-time must typically track an unbounded amount of information about past behaviour, even for relatively simple models, enforcing limits on precision due to the finite memory of the machine. However, quantum machines can require less information about the past than even their optimal classical counterparts to simulate the future of discrete-time processes, and we demonstrate that this advantage extends to the continuous-time regime. Moreover, we show that this reduction in the memory requirement can be unboundedly large, allowing for arbitrary precision even with a finite quantum memory. We provide a systematic method for finding superior quantum constructions, and a protocol for analogue simulation of continuous-time renewal processes with a quantum machine.

Our experience of the world manifests as a series of observations. The goal of science is to provide a consistent explanation for these, and further, make predictions about future observations. That is, science aims to provide a *model* of Nature, to describe the processes that give rise to the observations. It is possible to devise many different models that make identical predictions, and so it is desirable to have criteria that discern the ‘best’ model. One such guiding philosophy is Occam’s razor “plurality should not be posited without necessity”, which can be interpreted as requiring that a model should be the ‘simplest’ that accurately describes our observations.

This now leaves us with the question of how to determine the simplest model. The field of *computational mechanics* [1–3] seeks to answer this, defining the optimal predictive model of a process to be that which requires the least information about the past in order to predict the future, and uses this minimal memory requirement as a measure of complexity. There is motivation for preferring simpler models beyond the inclination for elegance; it allows one to make more fundamental statements about the processes themselves, due to their irreducible nature [3]. More pragmatically, it also facilitates the building of *simulators* (devices emulating the behaviour of the system [Fig. 1]) [4] for the process, as simpler models require fewer resources (here, internal memory).

Discrete-time processes have been well-studied within the computational mechanics framework [5–16]. However, it has recently been shown that quantum machines can be constructed that in general exhibit a lower complexity, and hence a lower memory requirement, than their optimal classical counterparts [17–24]. This substantiates the perhaps surprising notion that a quantum




FIG. 1: **Models and simulators.** (a) A system can be viewed as a black box which outputs what we observe, and our explanation of the observations forms a model. (b) We can build simulators to implement our models, in order to test them and make predictions about the future. Computational mechanics defines the optimal models to be those with the minimal internal memory requirement.

device can be more efficient than a classical system, even for the simulation of a purely classical stochastic process. This has recently been verified experimentally [25].

Continuous-time processes have also recently been the focus of computational mechanics studies [26, 27]. While the principles of optimality for discrete-time processes can be directly exported to the continuous-time case, systematic study of the underlying architecture has only been carried out for a restricted set of processes; renewal processes [27]. Renewal theory describes a generalisation of Poisson processes [28, 29], where a system emits at a time drawn from a probabilistic distribution, before returning to its initial state (one can also view this as a series of events separated by probabilistic dwell times). Despite their apparent simplicity, such processes have many applications, including models of lifetimes [30], queues

*Electronic address: physics@tjelliott.net

†Electronic address: milegu@quantumcomplexity.org

[31], and neural spike trains [13, 32, 33].

Here, we show that the quantum advantage can be extended to continuous-time processes. Focussing on renewal processes, we provide a systematic construction for determining quantum machines that require less information about the past for accurate future prediction than the optimal classical models. We provide a protocol that can be used to implement such quantum machines as analogue simulators of a renewal process. We then illustrate the quantum advantage with two examples. In particular, we find that while the classical machines typically need infinite memory, a quantum machine can in some cases require only finite memory. We then conclude by arguing that the finite memory requirement may be a typical property of quantum machines, suggest possible mediums for a physical realisation of the quantum machines, and discuss prospects for future work. A technical appendix is provided with additional details on the framework used, and the derivations for the examples.

Framework. A stochastic process \mathcal{P} is characterised by a sequence of observations (x_n, t_n) , drawn from a probability distribution $P(X_n, T_n)$ [34]. Here, the x_n , drawn from an alphabet \mathcal{A}_n , are the symbols emitted by the process, while the t_n record the times between emissions $n - 1$ and n . For shorthand, we denote the dual $\mathbf{x}_n = (x_n, t_n)$, and similarly \mathbf{X}_n for the associated stochastic variable. We denote a contiguous string of observations $\mathbf{x}_{l:m} = \mathbf{x}_l \mathbf{x}_{l+1} \dots \mathbf{x}_{m-1}$, and for a stationary process we mandate that $P(\mathbf{X}_{0:L}) = P(\mathbf{X}_{s:s+L}) \forall s, L \in \mathbb{Z}$.

We define the past of a process $\overleftarrow{\mathbf{x}} = \mathbf{x}_{-\infty:0} t_{0+}$, where 0 is the current emission step (i.e. the next emitted symbol will be x_0), and t_{0+} is the time since the last emission. Analogously, defining t_{0-} as the time to the next emission, we can denote the future $\overrightarrow{\mathbf{x}} = (x_0, t_{0-}) \mathbf{x}_{1:\infty}$ [27]. The causal states of the process are then defined according to a predictive equivalence relation [1]; two past sequences $\overleftarrow{\mathbf{x}}$ and $\overleftarrow{\mathbf{x}}'$ belong to the same causal state (i.e. $\overleftarrow{\mathbf{x}} \sim_e \overleftarrow{\mathbf{x}}'$) iff they satisfy

$$P(\overrightarrow{\mathbf{x}} | \overleftarrow{\mathbf{x}} = \overleftarrow{\mathbf{x}}) = P(\overrightarrow{\mathbf{x}} | \overleftarrow{\mathbf{x}} = \overleftarrow{\mathbf{x}}'). \quad (1)$$

The provably optimal (predictive) classical model that will faithfully reproduce the future statistics of the process operates on these causal states [1, 2]. This optimal model, termed the ‘ ε -machine’ of the process, is an edge-emitting hidden Markov model, where the edge emissions correspond to the symbols x_n . The quantity of interest for our study is the *statistical complexity* C_μ , which answers the question “What is the minimal information required about the past in order to accurately predict the future?”. It is defined as the Shannon entropy [35] of the steady state distribution π of the causal states S_j ;

$$C_\mu = - \sum_j \pi(S_j) \log_2(\pi(S_j)). \quad (2)$$

Because the state of an ε -machine is defined wholly and uniquely by its past, they satisfy the unifilarity condi-




FIG. 2: **Generative model for a renewal process.** Diagram depicting a generative model for a renewal process. The labelling indicates that a symbol 0 is emitted with probability 1, at time t with probability density $\phi(t)$, and returns to the same state.

tion: for a given initial causal state and subsequent emission(s), the current causal state is known with certainty [2].

As noted in the introduction, while large body of work has been devoted to studying the ε -machines of discrete-time processes, it is only recently that a similar causal architecture been uncovered for a restricted set of continuous-time processes, renewal processes [27]. Renewal processes form a special case of the above, where each emission occurs at an independent and identically distributed (IID) probabilistic time, and emits the same symbol. Such processes are defined entirely by this emission probability density $\phi(t)$, and the sequence is fully described by the emission times alone. It is useful to define the following quantities for a renewal process: the survival probability $\Phi(t) = \int_t^\infty \phi(t') dt'$; and the mean firing rate $\mu = (\int_0^\infty t \phi(t) dt)^{-1}$.

In Fig. 2 we show a generative model for such a process. Because of the IID nature of the process, the only relevant part of the past in predicting the future statistics is the time since the last emission, and this assists us only in predicting the time to the next emission [27]. Thus, the causal equivalence relation simplifies to

$$t_{0+} \sim_e t'_{0+} \Leftrightarrow P(T_{0-} | T_{0+} = t_{0+}) = P(T_{0-} | T_{0+} = t'_{0+}). \quad (3)$$

We label the causal states $S_{t_{0+}}$ according to the minimum t_{0+} belonging to the equivalence class. Depending on the form of $\phi(t)$, we can determine which t_{0+} belong to the same causal state. Notably, if $\phi(t)$ is Poissonian, the time since the last emission is irrelevant (as the decay rate is constant), and hence all t_{0+} belong to the same causal state - the process is memoryless and has $C_\mu = 0$. All other processes involve a continuum of causal states, which may either extend indefinitely, terminate in a single state at a certain time, or eventually enter a periodic continuum (see Appendix A). The steady state probability density $\pi(S_t)$ of the causal states depends on this causal architecture (Appendix B). We specifically highlight that states in the initial continuum have $\pi(S_t) = \mu \Phi(t)$; as we will later discuss, this is the only necessary part of the architecture once we turn to quantum causal states.

The statistical complexity of the process can be defined in correspondence with Eq. (2), by taking the continuous

limit of a discretised analogue of the process;

$$C_\mu = \lim_{\delta t \rightarrow 0} - \sum_{n=0}^{\infty} \pi(S_{n\delta t}) \delta t \log_2(\pi(S_{n\delta t}) \delta t). \quad (4)$$

This quantity will however either be zero (for a Poissonian emission probability density), or infinite (for all other distributions), due to the infinitesimal coarse-graining. Classically therefore, it is not the most enlightening measure of complexity, and has motivated earlier work on this topic [27] to instead consider use of the differential entropy for the statistical complexity; $C_\mu^{(\text{DE})} = - \int_0^\infty dt \pi(S_t) \log_2 \pi(S_t)$. While this quantity allows for a comparison of the complexity of two processes, we find it lacking as an absolute measure of complexity, as it requires one to take logarithms of dimensionful quantities, and loses the original physical motivation of being the information contained within the process about its past. Instead, we will employ the true continuum limit of the Shannon entropy Eq. (4) as the measure of a process' statistical complexity, accepting the infinities as faithfully representing that classical implementations of such models do indeed require infinite memory.

Quantum causal states. It has been shown that a quantum device simulating a discrete-time process can in general require less memory than the optimal classical model [17]. To do so, for each causal state S_j one constructs a corresponding quantum causal state $|S_j\rangle = \sum_{xk} \sqrt{T_{kj}^{(x)}} |x\rangle |k\rangle$, where $T_{kj}^{(x)}$ is the probability that a system in S_j will transition to S_k , while emitting symbol x . The machine then operates by mapping the state $|k\rangle$ with a blank ancilla to $|S_k\rangle$, following which measurement of the $|x\rangle$ subspace will produce symbol x with the correct probability, while leaving the remaining part of the system in $|S_k\rangle$. The internal steady state of the machine is given by $\rho = \sum_j \pi(S_j) |S_j\rangle \langle S_j|$. We refer to such constructions as q -machines, and their internal memory C_q can be described by the von Neumann entropy [35] of the steady state;

$$C_q = -\text{Tr}(\rho \log_2 \rho). \quad (5)$$

Unlike classical causal states, the overlaps $\langle S_j | S_k \rangle$ of different quantum causal states are in general non-zero, and hence $C_q \leq C_\mu$ (typically the inequality is strict); thence, the q -machine has a lower internal memory requirement than the corresponding ε -machine [17]. Physically, this memory saving can be understood as the lack of a need to store information that allows complete discrimination between two pasts when they have some overlap in their conditional futures. Note that while this quantum construction is superior to the optimal classical model, it is not known in general whether it provides the optimal quantum model. Indeed, for particular classes of process, constructions involving several symbol outputs are known that have even lower internal memory [19, 20].

We now seek to extend this quantum memory reduction advantage to the realm of continuous-time processes.

To do so, we first define a wavefunction $\psi(t) = \sqrt{\phi(t)}$. We can rephrase the survival probability and mean firing rate in terms of this wavefunction: $\Phi(t) = \int_t^\infty |\psi(t')|^2 dt'$; and $\mu = (\int_0^\infty t |\psi(t)|^2 dt)^{-1}$. Inspired by the quantum construction for discrete-time processes, we wish to construct quantum causal states $|S_t\rangle$ such that when a measurement is made of the state (in a predefined basis), it reports a value t' with probability (density) $P(T_{0-} = t' | T_{0+} = t)$. We may view the quantum causal state as a continuous alphabet (representing the value of t_{0-}) analogue of the discrete case, with only a single causal state (S_0) the system may transition to after emitting this symbol.

The probability density $P(T_{0-} = t' | T_{0+} = t)$ is given by $\phi(t+t') / \int_t^\infty \phi(t') dt' = \phi(t+t') / \Phi(t)$. By analogy with the discrete case we construct our quantum causal states as $|S_t\rangle = \int_0^\infty dt' \sqrt{P(T_{0-} = t' | T_{0+} = t)} |t'\rangle$, and thus:

$$|S_t\rangle = \frac{1}{\sqrt{\Phi(t)}} \int_0^\infty dt' \psi(t+t') |t'\rangle. \quad (6)$$

We emphasise that while the wavefunction is encoding information about time in the modelled process, the q -machine used for simulation may encode it in any practicable continuous variable, such as the position of a particle. The measurement basis used to obtain the correct statistics is of course that defined by $\{|t\rangle\}$.

When the first segment $[0, \tilde{t})$ of the continuous variable in a quantum causal state is swept across, if the system is not found to be in this region the state is modified by application of the projector $\Pi_{\tilde{t}} = \int_{\tilde{t}}^\infty dt |t\rangle \langle t|$ and appropriate renormalisation. When this projector is applied to the state $|S_t\rangle$, the resulting state is simply $|S_{t+\tilde{t}}\rangle$ displaced by \tilde{t} ; by correcting for this displacement the effect of the measurement sweep is exactly identical to the change in the internal memory of the machine if no emission is observed in a time period \tilde{t} , and thus the quantum causal states automatically update when measurement sweeps are used to simulate the progression of time.

The overlap of two quantum causal states can straightforwardly be calculated:

$$\langle S_a | S_b \rangle = \frac{1}{\sqrt{\Phi(a)\Phi(b)}} \int_0^\infty dt \psi(t+a) \psi(t+b). \quad (7)$$

By their very construction, these quantum states will automatically merge states with identical future statistics, even if we neglect the underlying causal architecture. Recall the causal equivalence relation Eq. (3). Since these probabilities wholly define the quantum states, if two quantum states have the same future statistics they are identical by definition. Due to the linearity of quantum mechanics, the steady state probabilities of the identical quantum states are added together to find the total probability for the state, much the same way as the underlying state probabilities are added together when merging states to form the classical causal states. Thus, when constructing the quantum 'causal' states, we are

at liberty to ignore the classical causal architecture as described in Appendix A, without any penalty to the information that is stored by the q -machine, and instead construct quantum states for all $t \geq 0$ according to the prescription of Eq. (6). Note that the causal architecture can still be used as a calculational aid.

Memory of continuous-time q -machines. From Eq. (7) we see that in general the overlap of the quantum causal states is non-zero, unlike the corresponding classical states, which are orthogonal. Because of this reduced distinguishability of the quantum causal states, the entropy of their steady state distribution is less than that of the classical causal states, and hence the amount of information that must be stored by the q -machine to accurately predict future statistics is less than that of the optimal classical machine, evincing a quantum advantage for the simulation of continuous-time stochastic processes. We will later show with our examples that this advantage can be unbounded, wherein q -machines have only a finite memory requirement for the simulation of processes for which the ε -machine requires an infinite amount of information about the past. Note that even when we consider coarse-graining the time since the last emission to a resolution of finite intervals δt we shall still see a quantum advantage due to the non-orthogonality of the quantum states.

The density matrix describing the internal state of the q -machine is given by $\rho = \int_0^\infty dt \pi(S_t) |S_t\rangle \langle S_t|$. As discussed above, we can construct the quantum states $|S_t\rangle$ for all t , in which case their steady state probability density $\pi(S_t)$ is given by $\mu \Phi(t)$. We thus find that the elements of the density matrix are given by $\rho(a, b) = \mu \int_0^\infty dt \psi(t+a) \psi(t+b)$. From this, we can construct a characteristic equation to find the eigenvalues λ_n that diagonalise the density matrix:

$$\mu \int_0^\infty db \int_0^\infty dt \psi(t+a) \psi(t+b) f_n(b) = \lambda_n f_n(a). \quad (8)$$

The information stored by the q -machine can then be expressed in terms of these eigenvalues; $C_q = -\sum_n \lambda_n \log_2 \lambda_n$.

We find that the memory requirement of q -machines is invariant under rescaling of the time variable in the emission probability density. Consider such a rescaling $t \rightarrow \alpha t = z$. The wavefunction scales $\psi(t) \rightarrow \sqrt{\alpha} \psi(z)$ (the factor in front due to normalisation), and the mean firing rate changes $\mu \rightarrow \alpha \mu$. Putting this into the characteristic equation Eq. (8), we find that the rescaled characteristic equation is of the same form, solved by the same eigenvalues λ_n and the rescaled (and renormalised) eigenfunctions $f_n^{[\alpha]}(z) = f_n(z/\alpha)$. As the memory stored depends only on the λ_n , it is hence unchanged: C_q is timescale invariant.

Building q -machine simulators of renewal processes. While we have explained in the abstract sense how one constructs the quantum causal states, it is interesting to also consider the structure of a device that




FIG. 3: **q -machine simulators of renewal processes.** (a) Analogue simulator for a discrete-time renewal process, where a continuous chain of qubits is used to encode the quantum causal state. The simulator sweeps along the chain and alters the future of the chain conditional on the current qubit, with the mappings $|0\rangle|1_n\rangle \rightarrow |0\rangle|1_n\rangle$ and $|1\rangle|0\rangle^{\otimes n} \rightarrow |1\rangle|\sigma_0\rangle$. Measurement of the qubit state signifies whether an emission occurs in a given timestep. (b) Analogue simulator for continuous-time renewal processes, where the quantum causal state is encoded into the position of a particle. The simulator sweeps along this position and generates additional particles encoding future emissions conditional on the presence of the particle. Detection of the particle signals an emission event.

would actually perform such simulations. In fact, a digital simulation of the process, that simply emits a sequence $t_{0-:L}$ on demand drawn from the correct probability distribution $P(T_{0-:L} = t_{0-:L} | T_{0+} = t_{0+})$ would be very straightforward to assemble in principle: one must prepare the state $|S_{t_{0+}}\rangle$, and $L-1$ copies of $|S_0\rangle$ (the states are all independent due to the renewal process emissions being IID). Measurement of the first state provides the t_{0-} , while measurement of the others provides the $t_{1:L}$. Because of the self-updating nature of the quantum causal states under partial measurement sweeps $[0, \hat{t}]$, measurement over such a range can be used to simulate the effect of waiting for a time \hat{t} for an emission.

However, this scheme is unsatisfactory as one must manually switch to a new state after each emission. Rather, a device that automatically begins operating on the state for the next emission after the previous state is finished would be preferable. We now describe such a construction, and even go a step further, by devising a setup that enables an analogue simulation of the process, and is thus able to provide emission times in (scaled) real time. For illustrative purposes, we first describe the protocol for discrete timesteps (that may be coarse-grained arbitrarily finely), and then discuss how it can be performed in continuous-time.

The procedure for the discrete-time case is as follows. Consider an infinite chain of qubits (two state quantum systems) labelled from 0 to ∞ . Using $|1_n\rangle$ to denote the state where all qubits are in state $|0\rangle$ apart from the n th, which is in state $|1\rangle$, we can express the discretised analogues $|\sigma_t\rangle$ of the quantum causal states $|S_t\rangle$ as $|\sigma_t\rangle = \sum_n \sqrt{P(T_{0-} = n\delta t | T_{0+} = t)\delta t} |1_n\rangle$, where

$P(T_{0-} = n\delta t | T_{0+} = t) \rightarrow \phi(t + n\delta t)/\Phi(t)$ as $\delta t \rightarrow 0$. The location n of the qubit in state $|1\rangle$ then represents the time $n\delta t$ at which the emission occurs. We initialise the system in state $|\sigma_{t_{0+}}\rangle$, according to the desired initial t_{0+} . The chain is then processed sequentially, one qubit at a time, by performing a control gate on the qubit, which has the effect of mapping the next block of the chain to the state $|\sigma_0\rangle$ if the qubit is in state $|1\rangle$, and doing nothing otherwise (explicitly, the mapping required is $|0\rangle|1_n\rangle \rightarrow |0\rangle|1_n\rangle \forall n \in \mathbb{Z}^+$ and $|1\rangle|0\rangle^{\otimes \infty} \rightarrow |1\rangle|\sigma_0\rangle$, where by construction these are the only possible input states). The qubit is then ejected from the machine (where measurement can be used to determine whether an emission event occurs at this time), and the machine then acts on the next qubit in the chain [Fig. 3(a)]. This operation has the effect of preparing the chain in a state that provides the correct conditional probabilities if no emission is observed, and prepares the state with the correct distribution for the next emission step if an emission is observed.

To operate this protocol in continuous-time, instead of encoding the state onto a discrete chain, we instead use a continuous degree of freedom, such as spatial position (henceforth referred to as the ‘tape’). As with the discrete case, we process sequentially along the tape, performing a unitary gate on the future of the tape, controlled on the current segment. Each emission step has its emission time encoded by the position of a particle on the tape [Fig. 3(b)]; the first particle on the tape is initialised in $|S_t\rangle = (1/\sqrt{\Phi(t)}) \int_0^\infty dx \psi(t+x)|x\rangle$, where x labels the position on the tape. Since the controlled unitary operation must be performed in discrete time, on a discrete length of tape, it is designed such that it acts, controlled on the presence of a particle in the block, by placing a particle in state $|S_0\rangle$, displaced to have its zero at the location of the control particle, and does nothing otherwise, akin to the discrete case above. Strictly, the gate should act in a nested fashion, by further generating an additional particle in an appropriately displaced state, when the new particle is placed within the current block. The machine then progresses to perform the same operation contiguously on the next block, while feeding out the previous block (equivalently, the tape can be fed through a static machine). Measurement of the positions of particles on the tape fed out then provides the simulated emission times.

Examples. We illustrate our proposal with two examples. We show for both these examples that not only is there a reduction in the memory requirement of the q -machine compared to the ε -machine, but also that the q -machine needs only a finite amount of memory, while the classical has infinite memory usage. Here we summarise the results, and the technical details may be found in Appendices C and D.

The first example is a uniform emission probability over the interval $[0, \tau]$. The corresponding emission probability density is $\phi(t) = 1/\tau$ for $0 \leq t < \tau$, and zero elsewhere [Fig. 4(a)]. The corresponding mean firing rate




FIG. 4: **Uniform emission probability.** (a) The corresponding emission probability density for a process with uniform emission probability in an interval $[0, \tau]$. (b) The classical memory C_μ required to simulate the process diverges logarithmically as the discretisation becomes finer (N states), while the quantum memory C_q converges on a finite value.

and survival probability are given by $\mu = 2/\tau$ and $\Phi(t) = 1 - t/\tau$ ($t < \tau$) respectively. The corresponding quantum causal states are given by $|S_t\rangle = \int_0^{\tau-t} dt' (1/\sqrt{\tau-t'})|t'\rangle$, and we can solve Eq. (8) to find that $\lambda_n = 8/(\pi(2n-1))^2$ for $n \in \mathbb{Z}^+$. We can use an integral test (see Appendix C) to show that $C_q = -\sum_{n=1}^\infty \lambda_n \log_2 \lambda_n$ is bounded, and moreover, that $C_q \approx 1.2809$. In Fig. 4(b) we show how the memory required by the q -machine tends towards this value as we use an increasingly fine coarse-graining of the discretised analogue of the process to approach the continuous limit, while the memory needed by the optimal classical machine diverges logarithmically.

For our second example, we consider a delayed Poisson process ($\phi(t) = (1/\tau_L) \exp(-(t - \tau_R)/\tau_L)$ for $t > \tau_R$ and 0 elsewhere), representing a process that exhibits an exponential decay with lifetime τ_L , and a rest period τ_R between emissions [Fig. 5(a)], forming, for example, a very crude model of a neuron firing. For this emission distribution we find that $\mu = (\tau_L + \tau_R)^{-1}$, and $\Phi(t) = 1$ for $t \leq \tau_R$ and $\exp(-(t - \tau_R)/\tau_L)$ for $t > \tau_R$. We can then show somewhat indirectly (see Appendix D) that the corresponding quantum memory requirement is bounded for finite τ_R/τ_L (and vanishes as this ratio tends to zero), while in contrast C_μ is infinite whenever this is non-zero. Further, due to the timescale invariance of the quantum memory, C_q depends only on this ratio, and not the individual values of τ_R and τ_L . Varying this ratio allows us to sweep between a simple Poisson process with lifetime τ_L , and a periodic process where the system is guaranteed to emit within an arbitrarily small interval at time




FIG. 5: **Delayed Poisson process.** (a) The corresponding emission probability density for a delayed Poisson process with rest period τ_R and lifetime τ_L . (b) Varying the ratio τ_R/τ_L sweeps between a (memoryless) Poisson process and a periodic emission process, with a corresponding increase in the quantum memory C_q required as the distribution becomes sharper. Here, C_q is calculated approximately from a very fine ($2^{14} + 1$ states) discretisation.

τ_R after the last emission. The quantum memory C_q correspondingly increases with this ratio as we interpolate between the two limits [Fig. 5(b)], with the pure Poisson process being memoryless, and a periodic process requiring increasing memory with the sharpness of the peak.

Discussion. We have shown that quantum devices can simulate models of continuous-time renewal processes with lower internal memory requirement than their corresponding optimal classical counterparts. Our examples evidence that this advantage can be arbitrarily large, compressing the need for an infinite classical memory into a finite quantum memory. Further, while we currently lack a proof, we suspect that this unbounded compression is a typical property of quantum machines. Our argument for this is as follows: Consider a discretised analogue of the process, with timesteps δt . When refining the discretisation to timesteps $\delta t/2$, we introduce an additional causal state for each one already existing. Classically, all these states are orthogonal, and the classical memory requirement increases, leading to the logarithmic divergence of C_μ as the timestep vanishes in the continuous limit. Contrarily, the overlap of adjacent quantum causal states $|S(t)\rangle$ and $|S(t + \delta t)\rangle$ is typically very large for small δt , and hence the additional interpolated states are very similar to the existing states. This overlap tends to one as the timestep vanishes, and thus as we tend to the continuum limit, the additional quantum causal states are essentially identical to those already

considered, and hence the entropy increase with refinement should vanish. Thus we expect that, with the exception of pathological cases (such as when the distribution $\phi(t)$ involves arbitrarily sharp peaks), the quantum memory requirement will be finite.

It is prudent of course, to remark on the experimental feasibility of our proposal. Recent works have succeeded in realising quantum machines for discrete-time processes, and demonstrating their advantage over classical devices [25]. Much effort has been spent on developing state-engineering protocols [36–40]. Ultracold atoms in optical lattices [41] may provide a route to realise discrete-time simulation, and of particular promise with regards to our proposal for simulating continuous-time processes, photon pulses have been shaped over distances of several hundred metres [42]. This would allow for the digital simulation we propose, with the requirement of an additional control element needed to achieve the ‘real-time’ analogue simulation. The on-going development of quantum technologies in a panoply of different systems holds much promise for the future implementation of our work.

Future theoretical work in this area can progress in many different directions, including the characterisation of other information-theoretic quantities [43, 44] for the continuous-time quantum machines, the application to study real-world stochastic systems, and the extension of the protocol to design quantum simulators for models of more general continuous-time processes. Further, while a quantum advantage over classical simulators has been demonstrated, the general optimal construction of quantum machines is unknown, and a subject for future investigation.

Acknowledgements. We thank Andrew Garner, Felix Binder and Suen Whei Yeap for useful discussions. This work was funded by Singapore National Research Foundation Fellowship NRF-NRFF2016-02. M.G. is also financially supported by the John Templeton Foundation Grant 53914 “Occam’s quantum mechanical razor: Can quantum theory admit the simplest understanding of reality?”. T.J.E. thanks the Centre for Quantum Technologies, National University of Singapore and the Department of Physics, University of Oxford for their hospitality.

Technical Appendix

A: Causal architecture for renewal processes.

For the purpose of determining their causal architecture, renewal processes can be classified as one of four types [27]. Here we review this classification, and describe the corresponding structure of their causal states.

The first class of processes are termed *eventually Δ -Poisson*. These are processes which exhibit a periodic structure under a Poissonian envelope at long times. Specifically, an eventually Δ -Poisson process is described by an emission probability of the form $\phi(t) = \phi(\tau +$




FIG. 6: **Causal architecture of renewal processes.** The causal architecture depends on the class of the renewal process, with (a) a single causal state for Poisson processes, (b) an infinite continuum of causal states for not eventually Δ -Poisson processes, and a hybrid of the two for (c) eventually Poisson and (d) eventually Δ -Poisson processes.

$(t - \tau) \bmod \Delta) \exp(-(t - \tau)/\tau_L)$ for times $t \geq \tau$, for some $\tau, \tau_L, \Delta \geq 0$. A special case of this, where $\Delta = 0$, forms our second class of process, *eventually Poisson*, while a yet further constrained form, with $\Delta = 0$ and $\tau = 0$, defines the familiar *Poisson* process class. Finally, the fourth class, *not eventually Δ -Poisson*, encompasses all other processes not of the above form.

We now present the causal architecture for each of the above processes [27]. Recall that two times since last emission t_{0+} and t'_{0+} belong to the same causal state iff they have identical conditional probability distributions $P(T_{0-}|T_{0+})$ for the time to the next emission t_{0-} . As a result of this, for a Poisson process all t_{0+} belong to a single causal state [Fig. 6(a)], as the decay rate is constant and independent of the time elapsed since the last emission.

For a not eventually Δ -Poisson process, it has been shown that no two different t_{0+} belong to the same causal state [27], and instead we have an infinite continuum of causal states, which the system traverses along between emissions, with all such continuum states returning to the same initial state immediately after an emission event (see Fig. 6(b)). The remaining two classes of process are hybrids of the previous two cases, with the eventually Poisson process beginning with a continuum that terminates in a single state once it reaches the Poisson stage at $t_{0+} = \tau$ [Fig. 6(c)], and the eventually Δ -Poisson process entering a periodic loop of continuum states of length Δ for times greater than τ (see Fig. 6(d)).

B: Statistical complexity of renewal processes. We now derive expressions for the statistical complexity

of the different classes of renewal process. We do so by considering a discretised analogue [16] of the continuum causal states, and take the limit of infinitesimal time intervals.

Starting with the case of a Poisson process, due to the single causal state we need not track any information about the time since the last emission; we thus have $C_\mu = 0$ and hence the process is memoryless. We next jump to the case of a not eventually Δ -Poisson process, where we have the perpetual continuum of states. Considering the discrete analogues σ_t of the causal states S_t , we have states at each $t = n\delta t$ for $n \in \mathbb{N}$, where δt is our discretised time interval. Consider now when the system is in causal state $\sigma_{n\delta t}$. In the next time interval δt the system will either make an emission, or progress along the continuum to $\sigma_{(n+1)\delta t}$. This latter event occurs with probability $\Phi((n+1)\delta t)/\Phi(n\delta t)$ (that is, the conditional probability that the system does not emit before $(n+1)\delta t$ given that it did not emit before $n\delta t$), and we can consider the system to fictitiously emit a \emptyset (null) symbol, representing the lack of a real emission. The alternative outcome is that the system actually emits a real symbol and returns to σ_0 , which occurs with probability $1 - \Phi((n+1)\delta t)/\Phi(n\delta t)$. This is illustrated in Fig. 7.

We can construct a transition matrix T that describes the evolution of the system for each timestep. This matrix evolves a state ω according to $\omega(t + \delta t) = T\omega(t)$. The transition matrix for a not eventually Δ -Poisson renewal process is given by

$$T = \begin{pmatrix} 1 - \frac{\Phi(\delta t)}{\Phi(0)} & 1 - \frac{\Phi(2\delta t)}{\Phi(\delta t)} & 1 - \frac{\Phi(3\delta t)}{\Phi(2\delta t)} & \dots \\ \frac{\Phi(\delta t)}{\Phi(0)} & 0 & 0 & \dots \\ 0 & \frac{\Phi(2\delta t)}{\Phi(\delta t)} & 0 & \dots \\ 0 & 0 & \frac{\Phi(3\delta t)}{\Phi(2\delta t)} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (9)$$

The steady state π is defined according to $\pi = T\pi$, and thus for $n \geq 1$ we have $\pi(\sigma_{n\delta t}) = (\Phi(n\delta t)/\Phi((n-1)\delta t))\pi(\sigma_{(n-1)\delta t})$, and hence iteratively we find $\pi(\sigma_{n\delta t}) = \Phi(n\delta t)\pi(\sigma_0)$. By integrating over $\Phi(t)$ to find the appropriate normalisation, we have that as $\delta t \rightarrow 0$, $\pi(\sigma_0) \rightarrow \mu\delta t$, and hence $\pi(\sigma_t) = \mu\Phi(t)\delta t$. Inserting this into the Shannon entropy Eq. (4), we have that the statistical complexity is given by

$$C_\mu = \lim_{\delta t \rightarrow 0} - \sum_{n=0}^{\infty} \mu\Phi(n\delta t)\delta t \log_2(\mu\Phi(n\delta t)\delta t), \quad (10)$$

which clearly diverges logarithmically as the argument of the logarithm vanishes (the number of terms in the sum grows linearly, while the prefactor to the logarithm decays linearly, thus effectively negating each other), and hence whenever such a continuum of causal states occurs (i.e. any renewal process that is not Poisson), the classical memory requirement is infinite. As noted in the main text, the self-assembly of the quantum causal states means that we can effectively treat any renewal process




FIG. 7: **Discretised analogue of a renewal process.** We can construct a discretised analogue of a renewal process, where at each time step δt the system can either emit a symbol 0 and return to the initial causal state, or emit nothing (signified by a null symbol \emptyset) and proceed to the next state in the chain. The continuous-time scenario follows as the $\delta t \rightarrow 0$ limit.

as not eventually Δ -Poisson in the quantum regime, and hence this steady state distribution is sufficient for our purposes. However, we will provide expressions for the steady state distributions and complexity of the eventually Poisson and eventually Δ -Poisson processes for completeness.

For eventually Poisson processes, the continuum has a finite length, after which the system resides in the final causal state until emission. The probability density for the continuum states is as for the above case, and the probability of occupation of the final state can be determined by considering the average time spent in this state, given by the lifetime of the state τ_L . Thus, the steady state occupation of this final state is $\pi(\sigma(\tau)) = \mu\Phi(\tau)\tau_L$, and the corresponding statistical complexity of the process is

$$C_\mu = \lim_{\delta t \rightarrow 0} - \sum_{n=0}^{N-1} \mu\Phi(n\delta t)\delta t \log_2(\mu\Phi(n\delta t)\delta t) - \mu\Phi(\tau)\tau_L \log_2(\mu\Phi(\tau)\tau_L), \quad (11)$$

where $N = \tau/\delta t$.

Finally, for eventually Δ -Poisson processes, we have two segments of continuum per emission; the initial line where each state is occupied at most once, and the periodic continuum where each state can be traversed multiple times per emission. There is a probability $\Phi(\tau)$ that the system reaches this periodic part on a given emission, and a probability $\Phi(\tau)\exp(-m\Delta/\tau_L)$ that it makes it through m circuits of this periodic component. Thus, the occupation probability in the steady state of the periodic continuum state σ_t ($\tau \leq t < \tau + \Delta$) is $\pi(\sigma_t) = \Phi(t) \sum_{m=0}^{\infty} \exp(-m\Delta/\tau_L) = \Phi(t)/(1 - \exp(-\Delta/\tau_L))$. This gives a statistical complexity of

$$C_\mu = \lim_{\delta t \rightarrow 0} - \sum_{n=0}^{N_1-1} \mu\Phi(n\delta t)\delta t \log_2(\mu\Phi(n\delta t)\delta t) - \sum_{n=0}^{N_2-1} \frac{\mu\Phi(\tau + n\delta t)\delta t}{1 - e^{-\Delta/\tau_L}} \log_2\left(\frac{\mu\Phi(\tau + n\delta t)\delta t}{1 - e^{-\Delta/\tau_L}}\right), \quad (12)$$

where $N_1 = \tau/\delta t$ and $N_2 = \Delta/\delta t$.

C: Technical details for uniform emission probability example. Here we provide details of the derivation of the boundedness of the quantum memory for the uniform emission probability case, along with the other associated quantities. Starting with the definition of the process

$$\phi(t) = \begin{cases} \frac{1}{\tau} & 0 \leq t < \tau \\ 0 & t \geq \tau \end{cases}, \quad (13)$$

we can straightforwardly obtain that $\mu^{-1} = \int_0^\tau (t/\tau) dt = \tau/2$, and

$$\Phi(t) = \int_t^\infty \phi(t') dt' = \begin{cases} 1 - \frac{t}{\tau} & 0 \leq t < \tau \\ 0 & t \geq \tau \end{cases}. \quad (14)$$

As the time between emissions is guaranteed to be less than τ , we can consider quantum causal states only within the interval $[0, \tau)$.

The form of the quantum causal states $|S_t\rangle = \int_0^{\tau-t} dt' (1/\sqrt{\tau-t})|t'\rangle$ follows directly from the definition Eq. (6), and we can construct the appropriate characteristic equation for the process:

$$\frac{2}{\tau} \int_0^\tau db \left(1 - \frac{\max(a, b)}{\tau}\right) f_n(b) = \lambda_n f_n(a). \quad (15)$$

Taking the second derivative of both sides of this equation, we find that

$$\lambda_n \frac{d^2 f_n}{dt^2} = -\frac{2}{\tau^2} f_n, \quad (16)$$

and hence the eigenfunctions are of the form $f_n(t) = A \exp(ik_n t) + B \exp(-ik_n t)$, with $k_n^2 = 2/(\lambda_n \tau^2)$. We must now determine the values of k_n that are valid, by substituting this solution into the original integral equation. Doing so, we find that for consistency, the following conditions must be satisfied: $A = B$; and $\cos(k_n \tau) = 0$. These are satisfied by $k_n \tau = (n - 1/2)\pi$ for $n \in \mathbb{Z}^+$ (zero and negative integer values of n produce solutions with the same eigenfunctions, by symmetry). Thus, we have that

$$\lambda_n = \frac{2}{(n - \frac{1}{2})^2 \pi^2} \quad n \in \mathbb{Z}^+. \quad (17)$$

We now wish to find the Shannon entropy of the λ_n . We first show that this entropy is bounded by using an integral test [45] for convergence, and then use this to provide a bounded range for C_q . Define the function $\zeta(n) = -\lambda(n) \log_2 \lambda(n)$, where $\lambda(n) = 2/(n - 1/2)^2 \pi^2$, the interpolated continuous analogue of the eigenvalues λ_n . The sum of $\zeta(n)$ over positive integers n gives the quantum memory cost, and we note that all such values of $\zeta(n)$ are finite. Further, we note that with the exception of $n = 1$, they satisfy $\zeta(n) > \zeta(n + 1)$ at these integer values, and hence the function is monotonically decreasing for $n > 2$ (specifically, the continuous function is decreasing for $n > 1/2 + \exp(\ln 2/2 + 2/\pi)/\pi$). Since the integral test requires the terms to be monotonically decreasing, we can sum up the terms to some finite $N > 1$, and then show the remainder of the terms converge.

Define the integral $I(N) = \int_N^\infty \zeta(x) dx$. We find that

$$I(N) = \frac{2 \left(\log_2 \left((N - \frac{1}{2})^2 \pi^2 \right) - 1 \right)}{\pi^2 (N - \frac{1}{2})} + \frac{4}{\pi (N - \frac{1}{2}) \ln 2}, \quad (18)$$

which is finite for all integer $N > 1$, and hence the sum converges, implying a finite value for C_q . We can also use this to bound the value of the sum from N to ∞ as being between $I(N)$ and $I(N) + \zeta(N)$. For $N = 2$, this allows us to bound $1.1046 \lesssim C_q \lesssim 1.4174$. By calculating additional terms in the sum prior to taking the bound, we can tighten this further; for $N = 10^6$, we find that $C_q \approx 1.2809$, with the additional neglected terms contributing $\mathcal{O}(10^{-5})$ as a correction.

D: Technical details for delayed Poisson process example. Here we provide the corresponding details of the derivations for the delayed Poisson process example. Recall that the process is defined by the emission probability density

$$\phi(t) = \begin{cases} 0 & 0 \leq t \leq \tau_R \\ \frac{1}{\tau_L} e^{-(t-\tau_R)/\tau_L} & t > \tau_R \end{cases}, \quad (19)$$

from which it is straightforward to calculate $\mu^{-1} = \tau_L + \tau_R$, and

$$\Phi(t) = \begin{cases} 1 & 0 \leq t \leq \tau_R \\ e^{-(t-\tau_R)/\tau_L} & t > \tau_R \end{cases}. \quad (20)$$

We can exploit the causal architecture discussed in Appendix A, and identify this as an eventually Poisson process. We thus have a continuum of causal states for $0 \leq t_{0+} < \tau_R$, and a single causal state for $t_{0+} \geq \tau_R$. From the mean firing rate and survival probability, we can see that the probability density for all continuum causal states in the steady state is given by $(\tau_L + \tau_R)^{-1}$, while the probability that the eventually Poisson causal state is occupied is $\tau_L/(\tau_L + \tau_R)$.

While we can approximately determine C_q by considering discretised time intervals $\delta t \ll \tau_L, \tau_R$ and see that the quantum memory appears to converge to a finite value




FIG. 8: Convergence of quantum memory for delayed Poisson process. (a) The classical memory requirement C_μ for the delayed Poisson process diverges logarithmically with finer discretisation ($N + 1$ states), while the quantum memory C_q appears to converge to a finite value. (b) Inspection of the eigenvalues of increasingly finer discretisation of the q -machine for the delayed Poisson process shows that the eigenvalues appear to fall off with a $1/n^2$ dependence. Plots shown for $\tau_R/\tau_L = 1$, and β is a normalisation constant chosen such that $\sum_{n=101}^\infty \beta/n^2 = \sum_{n=101}^{N+1} \lambda_n$ for the $N = 2^{15}$ case (eigenvalues ranked largest to smallest).

(see Fig. 8), it is not a simple task to find an analytical expression for the continuous time limit. Instead, we shall prove boundedness of the quantum memory by considering a less efficient encoding of the causal states, and proving that this suboptimal encoding scheme has a bounded C_q . Specifically, we encode the eventually Poisson state $|S_{\tau_R}\rangle$ to be orthogonal to the continuum states. The density matrix is now block-diagonal, with one block for the continuum states, and a single element for the eventually Poisson state, and thus the total entropy is the sum of the entropies of the two blocks. The eventually Poisson state block contributes a finite amount (as it is a single element), and we shall now show that the contribution from the continuum block is also finite.

We can use the characteristic equation Eq. (8) to find the entropy contribution from the continuum block. We find that the overlap of two quantum casual states is given by $\langle S_a | S_b \rangle = \exp(-|a - b|/2\tau_L)$, and hence

$$\frac{1}{\tau_L + \tau_R} \int_0^{\tau_R} db e^{-|a-b|/2\tau_L} f_n(b) = \lambda_n f_n(a). \quad (21)$$

Differentiating twice, we obtain

$$\frac{d^2 f_n}{dt^2} = -\frac{1}{4\tau_L} \left(\frac{4}{\lambda_n(\tau_L + \tau_R)} - \frac{1}{\tau_L} \right) f_n, \quad (22)$$

and so as with the previous case, the eigenfunctions are of the form $f_n(t) = A \exp(ik_n t) + B \exp(-ik_n t)$, now with

$$k_n^2 = \frac{1}{4\tau_L} \left(\frac{4}{\lambda_n(\tau_L + \tau_R)} - \frac{1}{\tau_L} \right). \quad (23)$$

Again, we substitute into the original integral equation Eq. (21), which results in the consistency equations $(1 - 2ik_n\tau_L)A = (1 + 2ik_n\tau_L)B$ and $\text{Im}((1 + 2ik_n\tau_L)^2 \exp(ik_n\tau_R)) = 0$. Thus, the valid k_n satisfy

$$\tan(k_n\tau_R) = \frac{4\tau_L k_n}{4\tau_L^2 k_n^2 - 1}, \quad (24)$$

which, with one exception has one solution in each interval $[m\pi, (m+1)\pi]$, $m \in \mathbb{N}$. The exception is during the interval in which $4\tau_L^2 k_n^2 = 1$, in which case there may be two solutions. For values of k_n for which $4\tau_L^2 k_n^2 \gg 1$, this is approximately satisfied by $k_n\tau_R = n\pi$, for $n \in \mathbb{Z}^+$, which leads to corresponding eigenvalues

$$\lambda_n = \frac{4}{\left(\frac{\tau_R}{\tau_L} + 1\right) \left(4n^2\pi^2 \frac{\tau_L^2}{\tau_R^2} + 1\right)}. \quad (25)$$

Strictly, these approximate eigenvalues are overestimations, as the solutions to Eq. (24) for large k_n are slightly larger than $n\pi/\tau_R$. However, as these $\lambda_n \ll 1$, this also overestimates their contribution to the entropy. We further note that when $4n^2\pi^2\tau_L^2/\tau_R^2 \gg 1$ (i.e. for sufficiently large n), the λ_n scale approximately as $1/n^2$. Again, this simplification overestimates the eigenvalues, and their contribution to the entropy. We can then break up the entropy into two parts; that from the finite number of terms preceeding the values of n for which the above approximations are valid (which, due to the finite number of terms, gives a finite contribution), and those that come from the terms in which we have such large n . These latter terms also have a finite contribution to the entropy due to their $1/n^2$ scaling, and hence the total entropy is finite. This completes our proof that the quantum memory requirement for the delayed Poisson process is finite, though unlike the previous example we do not have an analytical expression for this value.

-
- [1] J. P. Crutchfield and K. Young, Physical Review Letters **63**, 105 (1989).
 - [2] C. R. Shalizi and J. P. Crutchfield, Journal of Statistical Physics **104**, 817 (2001).
 - [3] J. P. Crutchfield, Nature Physics **8**, 17 (2012).
 - [4] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).
 - [5] J. P. Crutchfield and D. P. Feldman, Physical Review E **55**, R1239 (1997).
 - [6] P. Tino and M. Koteles, IEEE Transactions on Neural Networks **10**, 284 (1999).
 - [7] A. J. Palmer, C. W. Fairall, and W. A. Brewer, IEEE Transactions on Geoscience and Remote Sensing **38**, 2056 (2000).
 - [8] R. W. Clarke, M. P. Freeman, and N. W. Watkins, Physical Review E **67**, 016203 (2003).
 - [9] J. B. Park, J. W. Lee, J.-S. Yang, H.-H. Jo, and H.-T. Moon, Physica A: Statistical Mechanics and its Applications **379**, 179 (2007).
 - [10] C.-B. Li, H. Yang, and T. Komatsuzaki, Proceedings of the National Academy of Sciences **105**, 536 (2008).
 - [11] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney, Physical Review Letters **103**, 094101 (2009).
 - [12] W. Löhr, Entropy **11**, 385 (2009).
 - [13] R. Haslinger, K. L. Klinkner, and C. R. Shalizi, Neural Computation **22**, 121 (2010).
 - [14] D. Kelly, M. Dillingham, A. Hudson, and K. Wiesner, PloS one **7**, e29703 (2012).
 - [15] A. J. P. Garner, J. Thompson, V. Vedral, and M. Gu, arXiv preprint arXiv:1510.00010 (2015).
 - [16] S. E. Marzen and J. P. Crutchfield, Entropy **17**, 4891 (2015).
 - [17] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Nature Communications **3**, 762 (2012).
 - [18] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, arXiv preprint arXiv:1511.05738 (2015).
 - [19] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Scientific Reports **6** (2016).
 - [20] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Physical Review A **93**, 052317 (2016).
 - [21] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, Physics Letters A (2016).
 - [22] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, arXiv preprint arXiv:1609.03650 (2016).
 - [23] A. J. P. Garner, Q. Liu, J. Thompson, V. Vedral, and M. Gu, arXiv preprint arXiv:1609.04408 (2016).
 - [24] J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, npj Quantum Information **3**, 6 (2017).
 - [25] M. S. Palsson, M. Gu, J. Ho, H. M. Wiseman, and G. J. Pryde, Science Advances **3**, e1601302 (2017).
 - [26] P. M. Riechers and J. P. Crutchfield, arXiv preprint arXiv:1607.06526 (2016).
 - [27] S. E. Marzen and J. P. Crutchfield, arXiv preprint arXiv:1611.01099 (2016).
 - [28] W. L. Smith, Journal of the Royal Statistical Society. Series B (Methodological) pp. 243–302 (1958).
 - [29] V. S. Barbu and N. Limnios, *Semi-Markov chains and hidden semi-Markov models toward applications: their use in reliability and DNA analysis*, vol. 191 (Springer Science & Business Media, 2009).
 - [30] J. L. Doob, Transactions of the American Mathematical Society **63**, 422 (1948).
 - [31] V. V. Kalashnikov, *Mathematical methods in queueing theory*, vol. 271 (Springer Science & Business Media, 2013).
 - [32] W. Gerstner and W. M. Kistler, *Spiking neuron models: Single neurons, populations, plasticity* (Cambridge University Press, 2002).
 - [33] S. E. Marzen, M. R. DeWeese, and J. P. Crutchfield, Frontiers in Computational Neuroscience **9**, 105 (2015).
 - [34] A. Khintchine, Mathematische Annalen **109**, 604 (1934).
 - [35] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (2000).

- [36] F. Verstraete, M. M. Wolf, and J. I. Cirac, *Nature Physics* **5**, 633 (2009).
- [37] W. Yi, S. Diehl, A. J. Daley, and P. Zoller, *New Journal of Physics* **14**, 055002 (2012).
- [38] P. Hauke, R. J. Sewell, M. W. Mitchell, and M. Lewenstein, *Physical Review A* **87**, 021601 (2013).
- [39] M. K. Pedersen, J. J. W. H. Sørensen, M. C. Tichy, and J. F. Sherson, *New Journal of Physics* **16**, 113038 (2014).
- [40] T. J. Elliott, W. Kozłowski, S. F. Caballero-Benitez, and I. B. Mekhov, *Physical Review Letters* **114**, 113604 (2015).
- [41] M. Lewenstein, A. Sanpera, and V. Ahufinger, *Ultracold Atoms in Optical Lattices: Simulating quantum many-body systems* (Oxford University Press, 2012).
- [42] P. B. R. Nisbet-Jones, J. Dille, D. Ljunggren, and A. Kuhn, *New Journal of Physics* **13**, 103036 (2011).
- [43] R. G. James, C. J. Ellison, and J. P. Crutchfield, *Chaos* **21**, 037109 (2011).
- [44] S. Marzen and J. P. Crutchfield, *Entropy* **16**, 4713 (2014).
- [45] K. Knopp, *Infinite sequences and series* (Courier Corporation, 1956).